

UNITED STATES AIR FORCE RESEARCH LABORATORY

Cognitive Task Analysis and Work-Centered Support System Recommendations for a Deployed Network Operations Support Center (NOSC-D)

Terry Stanard
Marvin L. Thordsen
Michael J. McCloskey

Klein Associates Inc.
1750 Commerce Center Blvd. North
Fairborn, OH 45324

Patrick J. Vincent

TASC, Inc.
2555 University Boulevard
Fairborn, OH 45324

August 2001

Final Report for the Period October 2000 to August 2001

20021017 086

Approved for public release; distribution is unlimited.

Human Effectiveness Directorate
Deployment and Sustainment Division
Sustainment Logistics Branch
2698 G Street
Wright-Patterson AFB OH 45433-7604

NOTICES

When US Government drawings, specifications or other data are used for any purpose other than a definitely related Government procurement operation, the Government thereby incurs no responsibility nor any obligation whatsoever, and the fact that the Government may have formulated, furnished, or in any way supplied the said drawings, specifications or other data, is not to be regarded by implication or otherwise, as in any manner licensing the holder or any other person or corporation, or conveying any rights or permission to manufacture, use, or sell any patented invention that may in any way be related thereto.

Please do not request copies of this report from the Air Force Research Laboratory. Additional copies may be purchased from:

National Technical Information Service
5285 Port Royal Road
Springfield, VA 22161

Federal Government agencies registered with the Defense Technical Information Center should direct requests for copies of this report to:

Defense Technical Information Center
8725 John J. Kingman Rd., Ste 0944
Ft. Belvoir, VA 22060-6218

TECHNICAL REVIEW AND APPROVAL

AFRL-HE-WP-TR-2002-0165

This report has been reviewed by the Office of Public Affairs (PA) and is releasable to the National Technical Information Service (NTIS). At NTIS, it will be available to the general public, including foreign nations.

This technical report has been reviewed and is approved for publication.

FOR THE COMMANDER



MARK M. HOFFMAN
Deputy Chief
Deployment and Sustainment Division
Air Force Research Laboratory

REPORT DOCUMENTATION PAGE			Form Approved OMB No. 0704-0188	
Public reporting burden for this collection of information is estimated to average 1 hour per response, including the time for reviewing instructions, searching existing data sources, gathering and maintaining the data needed, and completing and reviewing the collection of information. Send comments regarding this burden estimate or any other aspect of this collection of information, including suggestions for reducing this burden, to Washington Headquarters Services, Directorate for Information Operations and Reports, 1215 Jefferson Davis Highway, Suite 1204, Arlington, VA 22202-4302, and to the Office of Management and Budget, Paperwork Reduction Project (0704-0188), Washington, DC 20503.				
1. AGENCY USE ONLY (Leave blank)		2. REPORT DATE August 2001		3. REPORT TYPE AND DATES COVERED Final - October 2000 - August 2001
4. TITLE AND SUBTITLE Cognitive Task Analysis and Work-Centered Support System Recommendations for a Deployed Network Operations Support Center (NOSC-D)			5. FUNDING NUMBERS C: F33615-99-D-6001 DO: 18 PE: 62022F PR: 1710 TA: D0 WU: 09	
6. AUTHOR(S) Terry Stanard, Marvin L. Thordsen, Michael J. McCloskey, Patrick J. Vincent				
7. PERFORMING ORGANIZATION NAME(S) AND ADDRESS(ES) Klein Associates Inc. 1750 Commerce Center Blvd. North Fairborn, OH 45324			8. PERFORMING ORGANIZATION REPORT NUMBER TASC, Inc. 2555 University Boulevard Fairborn, OH 45325	
9. SPONSORING/MONITORING AGENCY NAME(S) AND ADDRESS(ES) Air Force Research Laboratory, Human Effectiveness Directorate Deployment and Sustainment Division Air Force Materiel Command Sustainment Logistics Branch Wright-Patterson AFB OH 45433-7604			10. SPONSORING/MONITORING AGENCY REPORT NUMBER AFRL-HE-WP-TR-2002-0165	
11. SUPPLEMENTARY NOTES				
12a. DISTRIBUTION AVAILABILITY STATEMENT Approved for public release; distribution is unlimited.			12b. DISTRIBUTION CODE	
13. ABSTRACT (Maximum 200 words) This report presents the results of a preliminary Cognitive Task Analysis (CTA) of the deployed Network Operations Support Center (NOSC-D), and the implications of the NOSC-D current state of practice for the development of Work-Centered Support Systems (WCSS). The NOSC-D organization is a new addition to the Air Force (AF) enterprise network, and is tasked with the information assurance of a number of deployed Network Control Centers (NCC-D). Klein Associates studied the NOSC-D through a single data collection trip to the 12AF NOSC-D at Davis-Monthan AFB. During their trip, Klein Associates conducted Cognitive Task Analysis interviews with four (4) NOSC-D personnel. Because of the preliminary nature of the findings, the analysis is presented in conjunction with recommended future research objectives. This report is structured into three major sections. The first section addresses the responsibilities and challenges of the position within the NOSC-D. The second section addresses the coordination issues of the NOSC-D with other levels of the AF enterprise network. The third section suggests a detailed research plan for investigating the events and parameters of a distributed information attack that could be used in the table-top exercise.				
14. SUBJECT TERMS Work-Centered Support Systems Cognitive Task Analysis Information Warfare			15. NUMBER OF PAGES 44	
			16. PRICE CODE	
17. SECURITY CLASSIFICATION OF REPORT UNCLASSIFIED	18. SECURITY CLASSIFICATION OF THIS PAGE UNCLASSIFIED	19. SECURITY CLASSIFICATION OF ABSTRACT UNCLASSIFIED	20. LIMITATION OF ABSTRACT UL	

THIS PAGE LEFT INTENTIONALLY BLANK

PREFACE

This research was accomplished as part of the "Cognitive Task Analysis and Work-Centered Support System Recommendations for a Deployed Network Operations Support Center (NOSC-D)" task order under the Technology Readiness and Sustainment (TRS) contract (F33615-99-D-6001). The period of performance for this task order spanned from October 2000 through August 2001.

We would like to thank Capt. Scott Brown and Lt. Leigh Ottati of AFRL/HECA for their assistance in locating subject matter experts, and providing editorial comments on the draft report. We would also like to thank Capt. Dominguez of the 12AF NOSC-D at Davis-Monthan AFB for providing access to personnel for interview.

Table of Contents

INTRODUCTION	1
Work-Centered Support System	2
Cognitive Task Analysis	2
NOSC/NOSC-D FUNCTIONS AND CHALLENGES	2
The NOSC/NOSC-D Primary Functions/Positions	2
Network Defender	3
Enterprise Controller	3
Event Manager	4
Crew Commander (CC) and Crew Chief	4
Recognition-Primed Decision Model	5
Advanced Team Decision Making (ATDM)	7
Issues, Problems, and Challenges by Individual Positions	9
Network Defender	9
Enterprise Controller	11
Event Manager	12
Crew Commander	14
Issues, Problems, and Challenges Across Multiple Positions	17
Network Defender and Enterprise Controller	17
Network Defender and Event Manager	19
Network Defender and Crew Commander	20
Enterprise Controller and Event Manager	20
Enterprise Controller and Crew Commander	21
Event Manager and Crew Commander	21
AFCERT, NOSC/NOSC-D, IW FLIGHT, AND NCC: ASPECTS OF COORDINATION	22
Deployed Unit Attacks versus Air Force Wide Attacks	27
MODELING A DISTRIBUTED ATTACK	27
Detecting the Distributed Attack	28
Stages of a Coordinated Attack	29
A Distributed Attack Against Multiple Air Force Targets	32
SUMMARY	35
REFERENCES	37

List of Tables

Table 1. Recommended Research and Development Initiatives for NOSC/NOSC-D, AFCERT.	16
Table 2. Views into a Distributed Attack at Different Locations	34

List of Figures

Figure 1. The Recognition-Primed Decision (RPD) Model	5
Figure 2. ATDM: A framework for characterizing the elements of high and low-functioning teams	8
Figure 3. The research approach used to study the NOSC-D and recommend future WCSS research and development opportunities	9

List of Figures (continued)

Figure 4. NOSC & NOSC-D functional positions and relationships.....	10
Figure 5. Air Force network security interactions.	25
Figure 6. Wagon Wheel method.	26

THIS PAGE INTENTIONALLY LEFT BLANK

Cognitive Task Analysis and Work-Centered Support System Recommendations for a Deployed Network Operations Support Center (NOSC-D)

INTRODUCTION

This report presents the results of a preliminary Cognitive Task Analysis (CTA) of the deployed Network Operations Support Center (NOSC-D), and the implications of the NOSC-D current state of practice for the development of Work-Centered Support Systems (WCSS). The NOSC-D organization is a new addition to the Air Force (AF) enterprise network, and is tasked with the information assurance of a number of deployed Network Control Centers (NCC-D). Klein Associates studied the NOSC-D through a single data collection trip to the 12AF NOSC-D at Davis-Monthan AFB. This trip was done in conjunction with a data collection trip of the 12AF's information warfare flight (IWF) conducted by the Air Force Research Laboratory's Cyber Crew Interface Development (Cyber CID) program. During our trip, we conducted Cognitive Task Analysis interviews with four (4) NOSC-D personnel. Because of the preliminary nature of our findings, our analysis is presented in conjunction with recommended future research objectives. Where applicable, we have leveraged findings from previous research to enrich the analysis of the NOSC-D. This previous research includes a data collection trip to the Air Combat Command (ACC) (Langley AFB, VA) NOSC in support of the joint AFRL Information and Human Effectiveness Directorates' Air Force Enterprise Defense (AFED) program (Bradford, 2000a; Bradford, 2000b; Bradford, 2000c; Sweeney, 2000).

The report is structured into three major sections. The first section addresses the responsibilities and challenges of the positions within the NOSC-D. This includes analysis of single positions, and paired interactions of positions within the NOSC-D. To facilitate an understanding of the analysis, we outline a model of naturalistic decision making (NDM) and a framework for describing high and low functioning teams; both of which influenced our Cognitive Task Analysis approach. The analysis includes identification of opportunities where WCSS concepts and technologies could provide value added to the NOSC-D systems and displays.

The second section of the report addresses the coordination issues of the NOSC-D with other levels of the AF enterprise network. Our single data collection trip did not elicit a firm understanding of existing coordination processes, owing in large part to the recent activation of the 12AF NOSC-D to the AF enterprise, and the low level of NOSC-D experience of the personnel. However, previous research efforts have provided insights into a promising way of investigating this coordination using low-fidelity table-top exercises. The suggested exercise would mimic a large-scale distributed attack, a critical scenario that the AF enterprise has yet to encounter. We outline a detailed research plan for creating opportunities of interaction between layers of the AF enterprise, in order to provide opportunities to study them.

The third section of the report suggests a detailed research plan for investigating the events and parameters of a distributed information attack that could be used in the table-top exercise. Again, our research to date in the information warfare domain suggests that the AF enterprise has very limited experience responding to a distributed information attack. The current

state of readiness is poorly understood. Yet this scenario is one that is often cited as representative of true information warfare. In order to present a realistic table-top exercise of such an attack, a separate research effort may be necessary to study the parameters of it.

Work-Centered Support System

The Work-Centered Support System (WCSS) will be referenced throughout the report as a potential technology solution for task complexity issues confronting NOSC-D individual positions and their coordination. WCSS represents a framework for interface design that leverages advances in interface philosophy and artificial intelligence (Eggleston, Young, & Whitaker, 2000; Young, Eggleston, & Whitaker, 2000). A WCSS is an interface structured around the actual decisions and tasks of the domain, in this case, the information assurance activities of the NOSC-D. Middleware and intelligent agents retrieve relevant information that may be distributed across a variety of stove-piped systems, and fuse and present the information in a task-relevant format. Task complexity is reduced on the information retrieval and sense-making side through intelligent agents. Complexity is reduced on the interface side by fusing information from separate sources and then presenting the information in a format consistent with how the NOSC-D personnel think about and carry out their tasks. In this report, we primarily focus on higher-order decision making complexity issues that may benefit from the WCSS approach. Next we describe Cognitive Task Analysis, a means of performing a user-needs analysis for WCSS design.

Cognitive Task Analysis

Cognitive Task Analysis is a useful set of methods for uncovering the critical decisions, judgments and other cognitive requirements that the WCSS will ultimately support. Cognitive Task Analysis consists of tools for eliciting and representing general and specific knowledge in a domain. In essence, it represents the front-end user requirements analysis for the WCSS. The Cognitive Task Analysis goes beyond procedural knowledge and behavioral processes of tasks to get inside the head of the operator and tries to understand the "cognitive map" that guides decision making, problem solving, and other cognitive activities. Klein Associates has developed many specific knowledge elicitation tools for Cognitive Task Analysis. For the present study, we primarily used the Wagon Wheel method for capturing both team and individual cognitive requirements (Klinger, Phillips, & Thordsen, 2000). The goal of this method is to identify the main communication channels existing for each position on the team and the nature of the communications. We also adapted the Critical Decision Method which uses a challenging incident the performer has experienced in the past to probe for decision points, the cognitive strategies employed to overcome the challenges, critical cues that are important, and typical errors that are made (Hoffman, Crandall, & Shadbolt, 1998).

NOSC/NOSC-D FUNCTIONS AND CHALLENGES

The NOSC/NOSC-D Primary Functions/Positions.

There are four primary functions in the NOSC and in a NOSC-D: Network Defender, Enterprise Controller, Event Manager, and Crew Commander. In this section, we will discuss

these four functional positions and present an analysis of the challenges associated with their respective functions. We also make recommendations for future research and development to overcome the challenges. This includes recommendations for future Cognitive Task Analysis research, and the development of Work-Centered Support Systems as a technology solution.

We learned that the NOSC and NOSC-D perform the same core functions, but differ in level of staffing and theater of operation. The NOSC-D generally staffs one person per position, but the NOSC may staff multiple personnel in each position. The workload level of the NOSC is generally steady-state, but the NOSC-D can be overtaxed during deployment. During the initial phase of deployment, the NOSC-D must manage events associated with the enterprise set up. The NOSC-D is also more vulnerable to attack during early deployment, because it has not yet reached full and stabilized operability. We will reference these similarities and differences of the NOSC and NOSC-D during discussions of the core positions and the challenges they face individually and collectively. However, many discussions do not distinguish between the two organizations, because on many levels they operate in the same manner.

Below are brief descriptions of each of the four core positions within the NOSC-D. Additional information about individual positions and cognitive requirements within the NOSC (in particular) can be found within McCloskey and Chrenka, 2001.

Network Defender

The Network Defender (ND) monitors for intrusions. In other words, they detect and monitor "enemy" activity, rather than enterprise or friendly activity. They rely on intrusion detection software for this purpose (ASIMS: Automated Security Incident Management System). One of the primary purposes of this software is to generate alerts in response to possible intrusion attempts. These alerts supply the ND with key information that informs *why* the particular activity was flagged as a possible intrusion. The ND may then deepen their information search to discern attack patterns, the source of the intrusion, and the possible vulnerability the attacker may be trying to exploit. Ironically, the goal of the ND is *not* to track down the intruder, but to categorize the intrusion into one of 8 categories (0 thru 7): Normal Traffic (0), Unauthorized Root Access (1), Unauthorized User Access (2), Attempted Access (3), Denial of Service (4), Poor Security Practices (5), Probe (6), and Malicious Login / Virus (7). If there is a need to garner additional information, the ND will turn the incident over to analysts dedicated to formalizing the nature of the attack and the appropriate response to it. On occasion the vulnerability may be investigated using security auditing tools such as Internet Security Systems Scanner (ISS Scanner). In that case, the Network Defender will scan to the perimeter of the effected base network, but not perform internal network scans.

Enterprise Controller

The Enterprise Controller (EC), in contrast to the Network Defender, can be viewed as the "friendly forces" component of the NOSC/NOSC-D. The EC monitors and tracks the functional status of the enterprise. That is, the degree to which the enterprise is able to function as intended. At the NOSC level, this can include the systems of 10 or more *fixed* NCCs. In contrast, the NOSC-D enterprise is composed of *deployed* NCC-Ds. The EC monitors the

“health” of the enterprise network using a variety of customized, commercial off-the-shelf software such as Concord NetHealth and HP Open View. These tools are used to monitor and manage the status of services (e.g., web access), links, LAN, voice switches, and radio nets. Other parameters monitored include NCC server and router status, currency of software patches, and other concerns normally associated with network system administration.

Event Manager

The primary function of the Event Manager (EM) at both the NOSC and NOSC-D is to create and track event records using the commercial Remedy Trouble Ticketing System. We can envision situations where this position could be critical in providing the commander the information needed to understand the status of the enterprise information systems, the present weaknesses and vulnerabilities, and estimates regarding when different systems may be fully functional, or which are impaired for extended periods. Although the Crew Commander position actually briefs the AF commander about the overall system, it is the EM who must track any incidents or events within the enterprise. Currently, the EM uses the Trouble Ticketing System database software to create events and track their status. They also have access to some representational software so they can display the data in various graphical and tabular formats.

Crew Commander (CC) and Crew Chief

The Crew Commander (CC) is the command position. The CC position at the NOSC actually consists of two individuals for each shift. One is an officer and the other is a senior NCO. At the NOSC-D, the CC is staffed by one officer per shift. As a manager, the CC tracks how the NOSC/NOSC-D individual positions are functioning, tracks the team coordinating as a whole, and manages the professional development of team members. The CC is also the principal point-of-contact for the NOSC/NOSC-D. In this role, the CC acts as an officer in promoting or advancing the NOSC/NOSC-D concerns and issues regarding the enterprise, particularly with NCCs that the NOSC/NOSC-D serves. The CC also prepares and delivers daily standup briefs to AF officials regarding the enterprise. Therefore, the CC requires an understanding of the current issues faced by the other positions he manages. This includes the enterprise status, vulnerabilities and potential attacks, and significant events and their progress.

Before we progress to an analysis of the NOSC-D positions, it is important to note that many of the observations we report are influenced by our previous experiences working with teams and information technologies, and models of decision making and team performance in naturalistic environments. To facilitate discussion of our analysis results, we will briefly describe a model of naturalistic decision making, and a framework of team decision making and performance that differentiates advanced versus lesser performing teams. The significance of these models is that they guide our observational process, and also our thinking about the potential vulnerabilities and challenges of NOSC-D positions and team functions. These frameworks also guide our thinking about additional research needs, and development opportunities for WCSS.

Recognition-Primed Decision Model

The Recognition-Primed Decision (RPD) model (Figure 1) was developed through field studies examining the way experienced personnel actually make decisions (Klein, 1998). The model has been tested and has been supported by different research teams working in a great variety of settings including firefighter commanders confronting real-life fires (Calderwood, Crandall, & Klein, 1987), neonatal intensive care nurses (Crandall & Calderwood, 1989), and weather forecasters (Pliske, Klinger, Hutton, Crandall, Knight, & Klein, 1997).

The RPD model explains how people can use experience to react rapidly and make good decisions without engaging in costly option comparisons. The model contrasts with traditional Rational Choice models, which posit that decision makers generate multiple decision options, select multiple dimensions to evaluate each option, collect information to populate each

dimension for each option, determine relevant weights for each dimension, and finally, compute which decision option has the highest (or lowest) score across all dimensions and weights. The Rational Choice model is an appropriate way of modeling decision processes for decision makers having little or no experience in the decision making context, and when time and other resources permit extensive analysis. The RPD model explains the decision processes of experts operating in their field, who must make decisions with reduced and uncertain information, and under time pressure.

Figure 1 shows a composite view of the processes that operate within the RPD framework. To facilitate an understanding of these various processes, we will break them into manageable chunks that are color-coded. In routine situations (gray), the expert experiences the world in a dynamic context and comes to recognize a situation as typical. Situation typicality has four

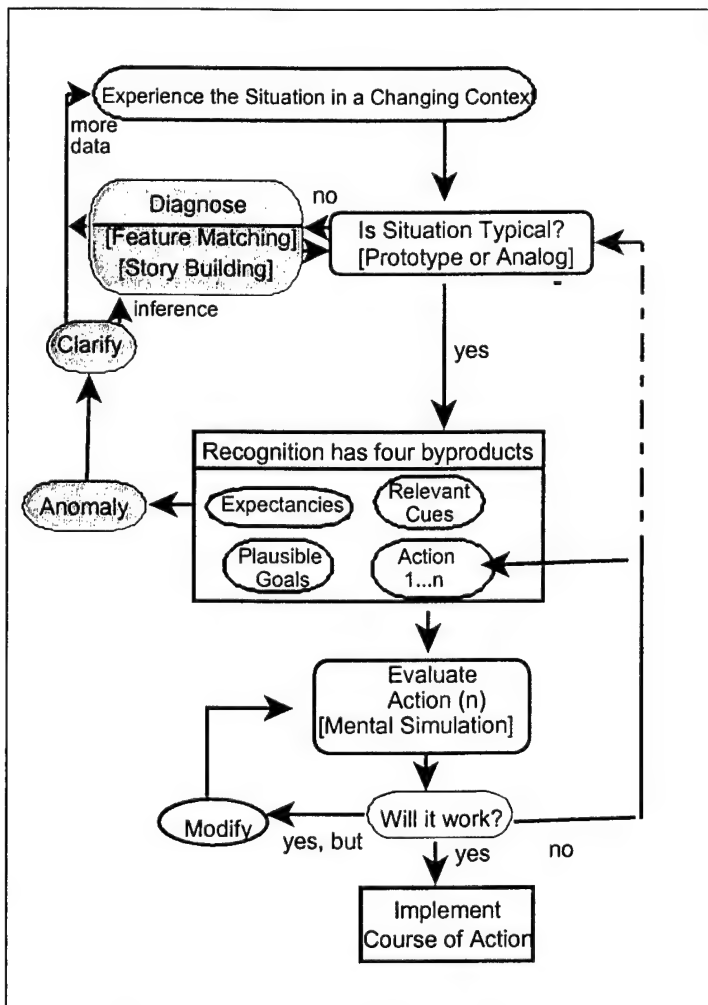


Figure 1. The Recognition-Primed Decision (RPD) Model.

components associated with it: relevant cues for that situation, expectancies, plausible goals, and plausible courses of action. Once the situation is recognized as familiar, a single course of action is obvious and is implemented. This basic process contrasts with rational-choice models of decision making, where several action alternatives are generated and the relative merits of each are evaluated before one is selected and executed.

Sometimes, the decision maker engages in more deliberate diagnosis because the situation is not immediately recognized, and information is actively sought to find cues and features that bring to light the nature or type of situation (blue). In other cases, the situation type may initially be recognized, but incorrectly. As events pan out, expectancies of future events are violated and the decision maker seeks clarification of the situation through the diagnostic process. Feature matching and storybuilding are two common strategies for diagnosing a situation.

Sometimes, a course of action has high potential costs associated with it. In such a case, the expert may mentally simulate and evaluate the course of action before execution (yellow). This may bring to light difficulties with the course of action, and may cause the expert to reject that course and consider another.

The RPD model is an example of a naturalistic decision-making (NDM) model. It describes how people actually think and process information under conditions of time pressure, ambiguous information, ill-defined goals, and changing conditions. The model focuses on experienced agents, working in complex, uncertain conditions, who face personal consequences for their actions. The model addresses situation awareness and problem solving as a part of the decision-making process. There are other models of naturalistic decision making that share similarities with RPD, but which contribute additional insights into decision-making processes. The reader may wish to consult other sources for a broader view of research and theory in naturalistic decision making such as Beach, 1993 (Image Theory); Pennington and Hastie, 1993 (story building); Montgomery, 1993 (Search for Dominance Structure); Rasmussen, 1993 (Decision Ladder).

Before the RPD model and other NDM models were developed, traditional decision researchers were aware that under certain task conditions people could not and would not use a Rational Choice strategy, but no one presented a coherent idea of what other strategies were available. The importance of the RPD model was to present a firm counter-example to Rational Choice. The significance of the RPD model is that it:

- § describes the most frequently used decision strategies,

- § explains how people can use experience to make difficult decisions, and

- § demonstrates that people can make effective decisions without going through a Rational Choice strategy.

The RPD model served as a template to direct our observations of the NOSC-D and focused our analysis on the NOSC-D functions and processes where decision making might be

most vulnerable. For instance, the RPD model posits that the decision maker “recognizes” a situation as typical. In the NOSC/NOSC-D realm, the Network Defender uses alert information to decipher the nature of an attack. The model shows that experienced decision makers can draw mistaken conclusions about what kind or type of situation exists when the information they are presented is misleading. Within the NOSC-D, the Network Defender may misapply a category to an alert, because the attacker has effectively clouded the type of attack in use. This is one of many such examples we will draw upon in our analysis.

Advanced Team Decision Making (ATDM)

The model of Advanced Team Decision Making (Figure 2) was developed by Klein Associates to help identify the key aspects of highly performing teams that set them apart from average or low performing teams (Zsombok, Klein, Kyne, & Klinger, 1992). Four key areas have been identified where the higher performing teams stand out.

- Team Resources concerns the team’s understanding of the skills and competencies of its members, their ability to handle basic procedures, and the team’s knowledge of other resources available to it. In other words, it addresses the team’s understanding of the “raw materials” that are available.
- Team Identity can be thought of as the “teamwork” portion of an advanced team’s repertoire. Does everyone know who does what? Is anyone being left out or not being utilized? Are people helping each other out when needed? Is anyone micromanaging?
- The Team Cognition component asks whether all the members of the team are headed for the same goals. Does everyone have the same picture? Are they behind the “power curve?” And, are they sidetracked or paralyzed by uncertainty?
- Finally, Team Metacognition refers to the team’s skill at monitoring its own performance and making appropriate adjustments. Successful team metacognition will show up as the team spotting and correcting problems in advance and by not being “time crunched.”

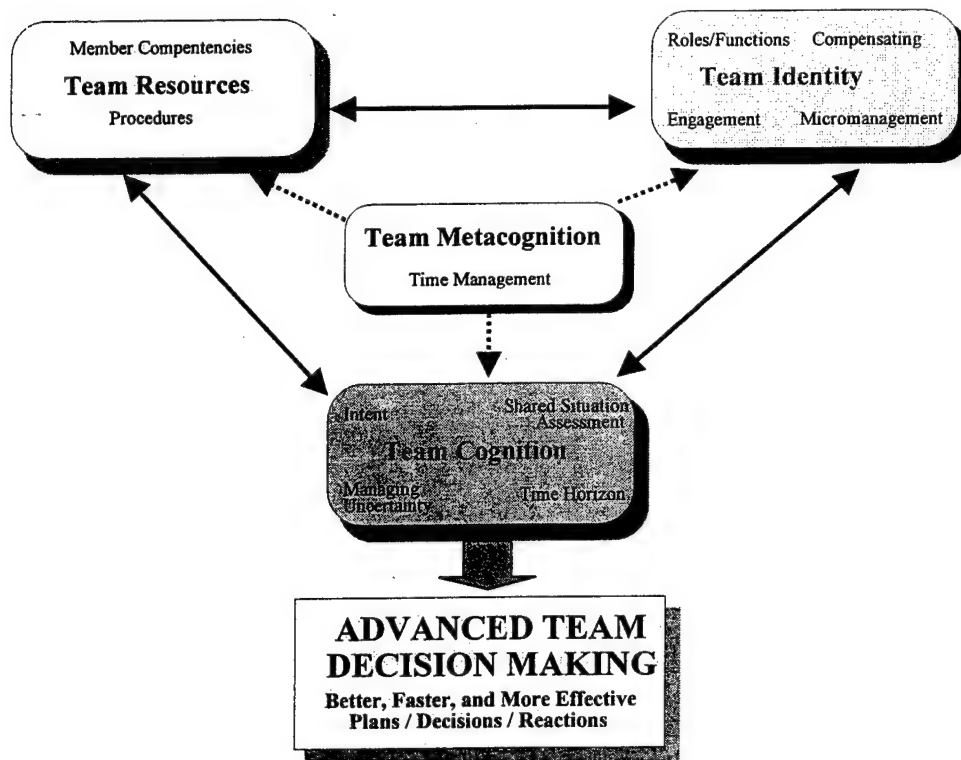


Figure 2. ATDM: A framework for characterizing the elements of high and low-functioning teams.

We can use ATDM to highlight critical areas for NOSC-D team performance and where they may be more likely to struggle. For instance, the Crew Commander may struggle to form a big picture view of the team function, because there is no display that supports the development of the big picture. A display for the Crew Commander could be designed that borrows information from existing displays that other positions already use, to build the big picture for the CC. The ATDM processes also suggest ways in which information technologies may be poorly designed. For instance, the CC display should do more than present all the detail already present in existing displays. It should be adapted to support the key decisions the CC has to make.

Figure 3 shows how we attempted to leverage our observations, previous research, the RPD model, and ATDM framework to help us identify the potential problems and challenges of the four NOSC-D positions. These resources help us articulate the problems and challenges the NOSC-D is likely to face, as well as the opportunities to leverage CTA as a research tool, and WCSS as a technology solution.

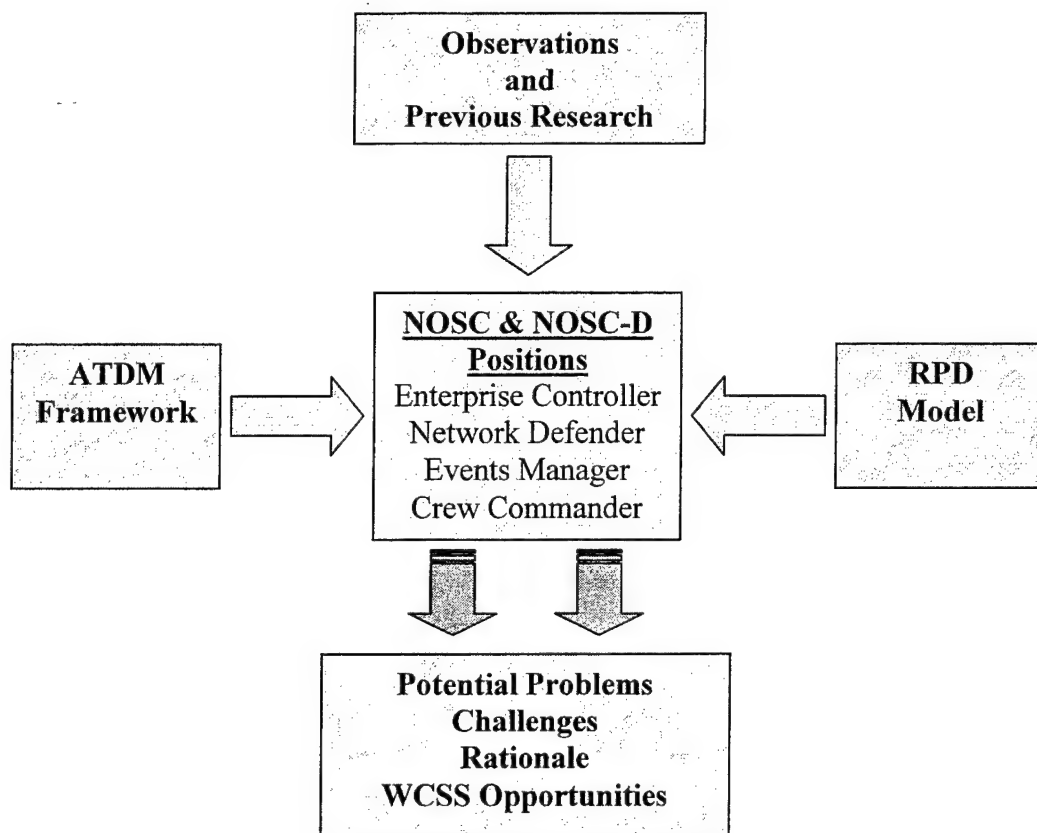



Figure 3. The research approach used to study the NOSC-D and recommend future WCSS research and development opportunities.

Issues, Problems, and Challenges by Individual Positions

Figure 4 presents a composite view of our analysis of individual and collaborative cognitive challenges associated with the NOSC-D. We will consult this framework as a “road map” during discussions about the challenges uncovered. The figure depicts the tools, major decisions and tasks of each position, and the relationships between each pair of positions. Opportunities for the development of WCSS are also summarized. These attributes for each position are next described in detail, followed by a review of each pair of positions.

Network Defender

The ND principally relies on ASIMS to deter, detect, isolate network intrusions, and recover compromised systems after attack. This software is pre-set to display “alerts” about known vulnerabilities based on known attack scripts. The attack scripts include port scans, or known vulnerabilities for particular pieces of software. We learned a relevant anecdote during



a previous data collection trip to the ACC NOSC. The ND told us about a particular program an NCC was using that had a vulnerability on a particular port. If an attacker knew of this bug, he/she could exploit it to gain access to the host computer. ASIMS was set up to display an alert anytime someone attempted to connect to the NCC host computer on that affected port address.

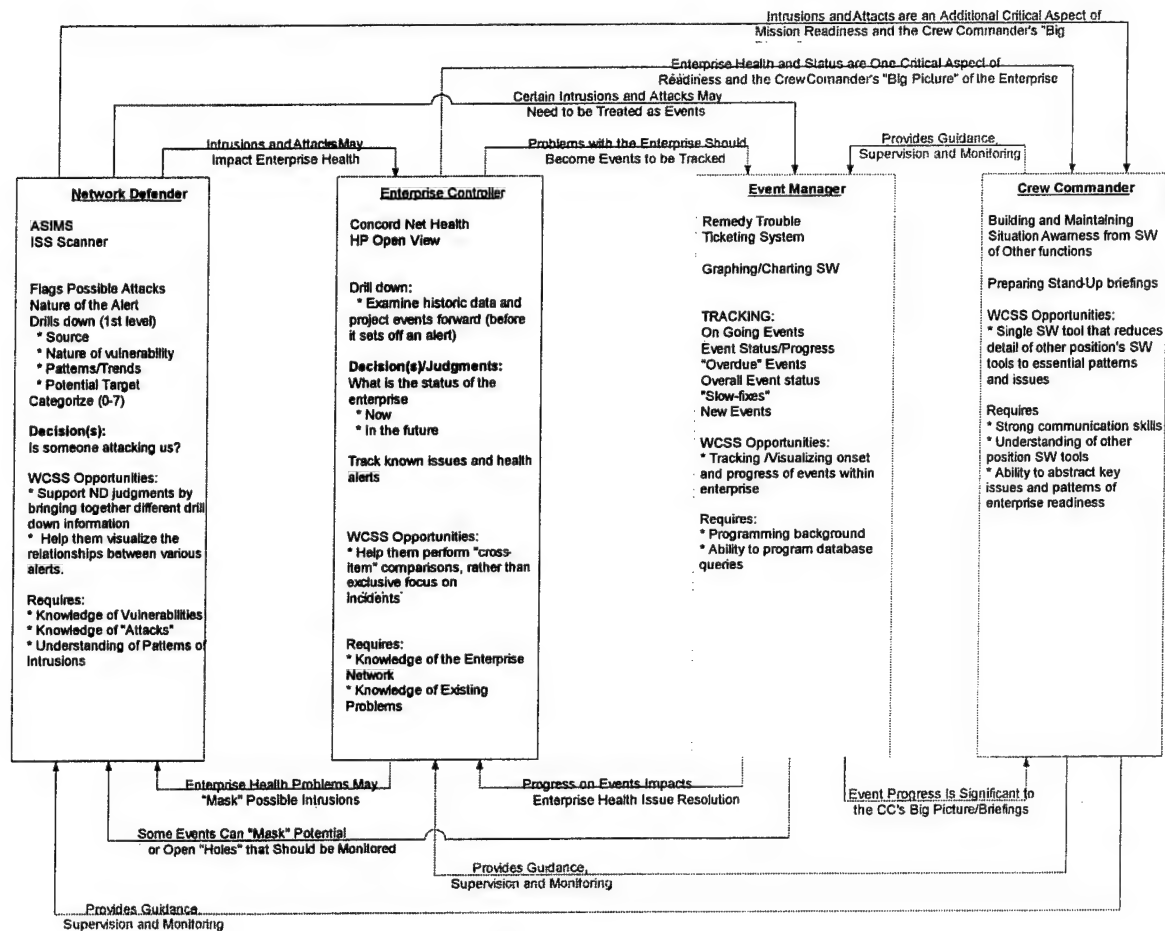


Figure 4. NOSC & NOSC-D functional positions and relationships.


The ND relies on the alerts to develop an understanding of potential attacks and intrusions. An alert may be investigated by using various Internet tools (e.g., "whois") to seek out additional information about a potential intruder. The information that can be gleaned from these tools includes the Internet service provider (ISP), particularly whether the Internet protocol (IP) address origin is within the military or private sector (e.g., .mil, .com, .gov). At any one time, the ND may be confronted with a large number of alerts. The sheer volume makes it difficult to discern patterns, particularly when the ND is forced to follow up the alert with additional analysis using Internet tools. Although individual alerts may represent individual, unrelated attacks, it is possible that the alerts point to similar or common attacks distributed across several originating IP addresses, against target ports, or over a period of time. The current displays can make these patterns difficult to recognize. The current alert formats only permit

simple pattern recognition, such as the number of times a potential intruder has tried to exploit a particular vulnerability, and other observed activity of that potential intruder.

ND Research and Development Opportunities. Because of the number of alerts per hour, the NDs usually just scan through them rapidly, identifying the known vulnerabilities, categorizing the alert, and then moving on. This makes it nearly impossible for them to devote dedicated resources to recognizing patterns across alerts. Doing so would require a deeper analysis and cross comparison of the information “behind” the alerts. The level of detail that they receive can interfere with them recognizing a larger picture, such as a distributed attack. Relating this to the RPD model, the ND may confront a certain volume of alerts that are normal or routine. But the current formatting of alerts can cloud the uncommon, anomalous patterns that underly individual alerts. This makes “recognition” difficult since there is so much chaff to sort through. A CTA opportunity would be to identify more specific details about the critical cues and *patterns* of cues that would indicate a real intrusion or attack versus benign patterns for a similar volume of alerts. The CTA would also focus on patterns that map onto each of the categories the ND normally uses. The results of the CTA could be used to design an alternate formatting of alerts that assists the ND in identifying patterns more readily. Agents would be designed, based on the CTA results, that actively seek out information about alerts that the ND reports as important, and present preliminary patterns across alert information parameters and a possible categorization scheme.

Enterprise Controller

The Enterprise Controller’s situation is different than that of the Network Defender. The EC uses sophisticated software that permits a quick, snap-shot view of the “health” of the enterprise network. The main display itself (Net Health) can provide views of network health at different levels of the enterprise. Usually, it is set to show the overall status of each NCC, with green, orange, and red lights indicating a healthy system, some problems, and disabled systems, respectively. But the EC can also “drill down” deeper within an NCC to examine the health of equipment in use. How deep they investigate depends upon how deeply the NCC network is configured—how many pieces of equipment they have set up to monitor.



However, we also learned from interviews with ECs that the current software and hardware they rely upon does not permit quick visualization of the impact of a single NCC problem on other NCCs, or the enterprise as a whole. In other words, potential relationships between the problems one NCC is having and problems other NCCs are experiencing are not immediately clear from current display symbology. In fact, the information presented in existing displays can instill a false sense of security or misunderstanding about enterprise health. The ability to identify and isolate individual problems can distract them from forming a larger overall picture of enterprise health.

Another function of the EC is to review logs and look for emerging patterns that may indicate a potential or impending problem. However, we have learned from previous work in the computer security domain that log files are as much a curse as a blessing (McCloskey, Stanard,


& Armstrong, 2001). It is easy to collect large volumes of information, and yet difficult to distill meaning from the information that is collected. Generally, log files are configured to feature too little information, or too much information, inhibiting the detection and recognition of patterns that signal impending enterprise health issues. Although log files can be configured to automatically flag known trends and patterns that indicate a problem, they do not support the diagnosis of unanticipated problems that lie beneath the surface features. Manual review of logs in search of critical health trends is a cognitively complex and time-consuming process that many ECs and other network administrators simply do not have the resources to support.

The issues of global enterprise health, and log file reviews are issues we have encountered before. Interviews with computer security personnel have shown that large variations exist in the level of expertise associated with security administration (McCloskey & Chrenka, 2001; McCloskey, Standard, & Armstrong, 2001). Some system administrators have developed better ways of working with available tools to develop the big picture of network health. Other system administrators are slaved to the automatic flags and warnings, and cannot see beyond the obvious indicators what is happening within their networked systems. They can also develop a false sense of security, if their entire assessment of attacks is based on the automated alerts. We have previously observed a very similar phenomenon with Air Force weather forecasters (Pliske et al., 1997). Some make excellent use of the wide array of computerized information sources to build an awareness of developing weather patterns. Others are slaved to one or a couple of information sources, and have little ability (or perhaps motivation) to seek out patterns, and build and test hypotheses about weather as it is forming.

EC Research and Development Opportunities. The expertise associated with developing a big picture of network health could be mined through Cognitive Task Analysis. Cognitive Task Analysis interviews could be conducted with ECs, but also non-military system administrators tasked with similar enterprise health management issues. Experts can be interviewed about critical events where they were confronted with singular network status problems that had implications for the entire enterprise, or at least other links within it. These incidents can be mined to construct a timeline of the event, identify actual decisions and judgments, cues that emerged with diagnostic value, and strategies that were used to make sense of the cues. The decision strategies, critical cues and patterns of cues could then be transferred to a) design better interfaces that make developing network health issues more apparent; b) create intelligent agents that are programmed to better recognize the same emergent health issues as experts; or c) develop training programs that instruct other ECs in the tactics the experts use to identify and manage developing network health problems. Administrators could be interviewed about their specific use of Net Health, and also log files, since both are industry standard tools. The selection of commercial off-the-shelf software by the NOSC/NOSC-D is advantageous, since it allows comparisons of Air Force strategies with the strategies of non-military administrators, many of whom may be more experienced than AF personnel.

Event Manager

Compared to other NOSC positions, the Event Manager (EM) position appears relatively undemanding. The key function of the EM is to log any new enterprise events that are submitted from within the NOSC or from



an NCC, and to track their progress. For example, if an NCC has a bridge router go down, it will likely become an event that the EM logs and tracks. The EM will enter it in the database, collect all relevant information (when, what, where, what is being done, who is remedying it, etc.), and then ensure this information is included in the daily summary materials provided to the Crew Commander for reports and briefings. In the NOSC, these EM tasks are relatively routine because they operate within a relatively "steady state" environment. The real challenges surface when we consider the EM at a NOSC-D, specifically during initial deployment, when all systems are not yet fully-configured and running. However, before we examine the deployment scenario closer, we will discuss the challenges associated with the normal, steady-state EM operations.

Although the position is not currently demanding, the information that the EM tracks is actually quite critical. It is essentially the list of "trouble tickets," and provides an overview of all the significant problems, what is being done about them, when they *should* be completed, which ones are behind schedule, and usually some information about who is responsible for the fix. This material is entered into a standard database and the EM can query it for different combinations of problems, etc. For example, they can sort the data by "routers," or "NCCs." While helpful, the creation of database management queries requires at least a working knowledge of the database architecture. Ironically, every EM we interviewed said they had no need for any additional software or tools. This was initially confusing, because database querying is not a "drag-and-drop," or simple procedure. Their stance became more clear when we realized that almost all of EM personnel have a programming background and enjoy writing queries. They do not want any additional technologies. It is like the person who refuses to move to a Windows desktop operating system because they like the control that command-line interfaces supply (such as Unix). Thus, the EM position may require less expertise, as long as programmers consistently occupy the slot. However, this may not always be the case. The trouble ticketing system is a skill to be learned, and not all will learn it equally well if they do not have a programming background.

We also learned from interviews with EM personnel that the database capabilities are many. It can provide "sorts" of events on many or few dimensions. However, it may not help them recognize important patterns in events. For example, it may be important to see all the events attributable to one NCC, or all events that affect a particular router. They may not be aware of what events are on schedule, and which are behind schedule. Such information may or may not be critical to the EM, but it is very important for the Commander, who needs to build a big picture about trends to properly brief his commander on enterprise status.


EM Research and Development Opportunities. The EM needs tools that represent the event data in a format more relevant to the needs of the Commander, as well as other NOSC and NOSC-D positions. Cognitive Task Analysis can be used to elicit the information needs of NOSC and NOSC-D positions regarding events and event tracking. The results of the Cognitive Task Analysis can be used to drive the features of a WCSS that present the critical event tracking information. For instance, the WCSS could present all the operative events in a single glance, with a common timeline feature that shows which are new, which have remained unaddressed (and for how long), and which events are progressing quickly and perhaps require less attention. A WCSS was created for the airlift planning domain that served a similar function. The WCSS

allowed dispatchers to better track which tankers experienced scheduling conflicts. It displayed flights with overlapping timelines, so that the conflicts were more apparent.

The ramp-up during initial deployment of a NOSC-D is a specific scenario with important surrounding research issues. During ramp-up, the EM is likely to be inundated with a significant number of incidents. Each NCC-D will set up and configure their network, with equipment recently shipped across many miles. There are going to be problems during installation. This will generate a large number of events. So although tracking events in steady-state may be manageable, during ramp-up periods it will be much more challenging. The EM at the NOSC-D is staffed by one individual per shift. This EM may be overtaxed logging and tracking events using a standard trouble ticketing system. A WCSS tool that permits them to examine unfolding patterns in events will serve the needs of the EM, as well as the Crew Commander who maintains the big picture view of enterprise to pass along to the AF commander during briefings.

Crew Commander

The role of the CC is to oversee the operation of the NOSC/NOSC-D; make sure that the other positions are operating at adequate levels; provide officer-level “muscle” when needed to ensure changes occur (e.g., fixes on routers, installation of security patches, etc.); brief the enterprise status to command personnel; and take ultimate responsibility for the NOSC/NOSC-D level of the overall enterprise. To monitor the NOSC/NOSC-D status, the CC primarily “eavesdrops” using the tools that are provided for the other three functions (ND, EC and EM). In addition, they rely heavily on email and telephone calls.



One problem the CC encounters is reliance on software tools that are better tailored for the positions for which they were designed (i.e., ND, EC, and EM). The formatting of the information available from these displays is often too detailed for the CC. For example, the EM needs to track individual events, so they require very specific information about each event. The event ticket features the name of the supplier point-of-contact who is coordinating the shipment of a new router to an NCC. In contrast, what the CC needs to know is that there is a router problem, it is being handled by having a new one shipped, and it is expected to be completed by X date. This higher-order information is more difficult to distill from the detail the EM normally is presented.

During the development and testing of the Advanced Team Decision Making model, we found that micromanagement was often the by-product of having too much detailed information thrust upon the manager. It was not just a function of the manager’s work style. For example, if the supervisor is continuously fed extremely detailed information, the details can pull them down “into the weeds.” By providing detailed information, the manager often ends up working at that level of detail. So, having the CC work with the level of information detail available to them through the ND, EC, and EM tools may inadvertently set the stage for micromanagement.

CC Research and Development Opportunities. The principles of WCSS can be leveraged to create a display for the CC that presents a summary view of information gleaned from ND, EC, and EM displays. A front-end Cognitive Task Analysis could be conducted to learn what information is relevant from each of the other position's hardware and software, and how that information is utilized by the CC. Intelligent agents are then created that search the networked tools that each position uses, and present that information in a format commensurate with their situation assessment activities of the CC.

Table 1 summarizes the Cognitive Task Analysis and WCSS research recommendations outlined above for the cognitive challenges confronting the individual positions at the NOSC/NOSC-D and the Air Force Computer Emergency Response Team (AFCERT). It also presents cognitive challenges and research objectives for relationships between pairs of positions, which are discussed next. This research agenda is preliminary, and requires the review of subject-matter experts more fully acquainted with the current standing of the NOSC and NOSC-D.

Table 1.

Recommended Research and Development Initiatives for NOSC/NOSC-D, AFCERT

POSITION	ISSUE	Uncovered by CTA	WCSS Implications
Network Defender (ND)	Recognizing alert patterns	Relevant cues for intrusion pattern recognition	Specialized agents retrieve cue information, visually formatted to support pattern recognition
Enterprise Controller (EC)	Understanding impact of NCC problem on other NCCs; Managing voluminous logs; Improving early problem detection	Expertise in building big picture of enterprise health	Agents collect emerging health issue data; Visualization to support correlation of NCC health issues
Events Manager (EM)	Early deployment of NOSC-D - recognizing patterns of events - tracking event progress	Event patterns and progress information	Visually represents event progress and event interrelationships
Crew Commander (CC) ND, EM, EC - CC	Interpreting detailed displays used by other NOSC/NOSC-D positions	Critical information needs, and appropriate information level to build big picture of enterprise status	Suite of tools providing overview of intrusions, health issues, and events
ND-EC	Deployed NOSC-D - Understanding impact of attacks on enterprise health - Understanding impact of enterprise health on attack vulnerability	ND decisions and information needs; EC decisions and information needs; Information availability for each position to satisfy the other	Agents mine data sources from ND tools for EC use, EC tools for ND use
EC-EM	Verifying correspondence of network health issues and logged (EM) events	EC decisions and information needs; EM decisions and information needs; Information availability for each position to satisfy the other	Highlight discrepancies between network health issues and events
AFCERT	Managing intrusion alert (and enterprise health issue?) overload; Identifying alert patterns across NCCs, NCC-Ds, NOSCs, NOSC-Ds	Alert patterns across enterprise levels	Alert display that better supports pattern recognition across enterprise levels


Issues, Problems, and Challenges Across Multiple Positions

In the preceding section, we addressed the general functions and challenges of each NOSC/NOSC-D position independently. However, referring back to Figure 4, there are additional opportunities when we examine the *relationships and interactions* between each of the positions. In this section we discuss the relationships, vulnerabilities, and opportunities for each of the six position dyads: ND-EC, ND-EM, ND-CC, EC-EM, EC-CC, and EM-CC. In discussing each position pair, we will explore issues such as:

- How can information at one position support the other position?
- What are the implications of sharing/not sharing this information?
- What insights from naturalistic models of decision making and teams can we draw for the position pair?
- Are there any potential opportunities for WCSS?

Network Defender and Enterprise Controller

The Enterprise Controller and Network Defender have an interesting relationship. A rough analogy would be the Operations (J3) and Intelligence (J2) cells of a joint air operations center. The EC monitors the “friendlies” and the ND tracks the “enemies.” Operations and Intelligence independently contribute information that has an impact on the other. Likewise, the enemy status that the ND tracks can have a direct impact on friendly status that the EC tracks, and vice versa. Friendly status (NCC placement, operational equipment, mission readiness, etc.) can impact where those NCCs may be vulnerable to enemy attacks. Friendly activities can occasionally be confused with and occasionally “mask” enemy activities, especially if the friendly activities draw your attention away from other critical areas. Thus, the EC may classify an event as benign, attributing it to normal equipment failure, when in fact, it represents an attack.



However, during our interviews at the ACC NOSC, we learned that the ND did not feel they had any information that the EC really needed. The EC also believed they had no information of use to the ND. But both the EC and the ND said they could use information that the other had. Referring back at the RPD model and ATDM framework, we can note several things. First, both positions have a low level of understanding of the critical information they possess that could be valuable to the other position. This may be attributed to an incomplete understanding of the complete roles and functions of the other stations (“Team Identity” in Figure 2). Because of this misunderstanding and the resultant lack of information sharing, both positions are deprived of some critical information that could help them better recognize and sort out real events from “chaff.”

Second, it appears that the EC does not understand that certain enterprise events may assist the ND in detecting some attacks or intrusions. In return, the ND does not understand that the EC needs to be advised of certain types of attacks and intrusions, so that the EC can anticipate the impact on the enterprise health. Finally, neither position understands that they each can act as an "advanced scout" for the other. For example, if the ND suspects the appearance of a certain kind of attack, the ND could have EC closely monitor specific parts of the enterprise to help them confirm or deny the ND's suspicions. In turn, if certain enterprise problems arise, the EC may want to inform the ND to watch these areas closely because additional vulnerabilities may have temporarily appeared.

ND-EC Research and Development Opportunities. These interactions all refer to the various status, pattern, and trend data that could be usefully shared between positions. This suggests a research and development opportunity. The ND and EC positions function relatively independently and work less collaboratively than they could, or should. Cognitive Task Analysis tools can be used to identify the critical decisions and judgments the ND and EC make independently, and the critical information that other position can offer in support of these judgments. A WCSS could be created with agents programmed to sift the data sources the ND uses, that the EC could use, and the data sources the EC uses that the ND could use. This cross-relevant information could then be presented in concert with information already presented to the ND and EC, in a format that supports their individual decision requirements.

Note that the cross-relevancy of information between the ND and EC is particularly critical for the deployed NOSC-D, especially during the "ramp-up" period. Several factors are present during the initial phase of deployment. First, a crisis has probably arisen that prompts the need for deployment. We are usually deployed to *help* one group or nation, while another believes we are sided *against* them, and may try to attack us. This puts the NOSC-D in harms way to a degree that the in-garrison NOSC will not typically experience. Second, during initial deployment, the enterprise will be more vulnerable. Not all of the equipment will be working and/or configured properly. There will be more holes, and thus, more vulnerabilities. So we have an enemy, and are vulnerable. Third, as the systems are coming on-line, there will be a significant amount of "noise" due to the equipment malfunctions and all the other factors associated with bringing a system on-line. This noise means it will be harder to sort through all of the events to be able to see any real threats that may slip (or sneak) through. Finally, the attention of deployed force will be focused (regardless of noise) on bringing their own systems on-line. Our forces may be less ready in the attentional aspects until the systems are up and running.

We can envision that a NOSC-D deployment could be a golden opportunity for hackers, cyber-terrorists, or other cyber-warriors. While some of the issues we have raised concerning ND-EC relations may not seem critical during steady-state NOSC operations, the relationship may represent an Achilles' heel at the NOSC-D.

Implications for AFCERT. Our previous study of AFCERT analysts revealed an important issue that we also see present in the ND-EC relationship. The NOSC-D and NOSC monitor all NCCs within their respective enterprise. But AFCERT monitors all activity across the Air Force networks. If an average NOSC monitors and classifies 50 alerts per hour, there are


50 x 8 alerts delivered to AFCERT at any one time (there are 8 NOSCs). The Air Force Network Operations Center (AFNOC) will monitor information about enterprise health, comparable to the EC function (although we do not have direct evidence to support or deny this). Thus, like the ND and EC, AFCERT and AFNOC analysts may have an even greater need to see the implications of intrusion alerts for Air Force network health, and the influence of various aspects of network health on vulnerability to attack.

However, the AFCERT vantage point gives them a more potentially improved view of attack patterns than any NOSC or NOSC-D has. If a NOSC-D is experiencing increased attack, this should show in the relative level of alerts the NOSC-D encounters, compared with other ingarrison NOSC-Ds. Similarly, the NOSC-D can be better informed by AFCERT about the nature of current alert patterns. The NOSC-D may falsely attribute the pattern to a concerted attack, but in fact, the alert level has not risen above the norm for Air Force networks overall.

Thus, AFCERT may actually benefit more from a Work-Centered Support System that displays patterns of alerts across NOSC/NOSC-Ds than would an individual NOSC/NOSC-D. A WCSS created for purposes of information sharing between the ND and EC position would probably be all the more significant for AFCERT and AFNOC, if those organizations monitor both intrusion alerts and enterprise health issues concurrently.

Network Defender and Event Manager.

The relationship between the ND and the EM is not as strong as between the ND and the EC. However, there are still areas where information exchange between the two would be beneficial. Specifically, the EM may log events that have implications for the focus of the ND efforts and alertness. For example, a previously unknown vulnerability is discovered on a particular type of router and a "patch" is developed that will have to be installed on every one of these routers in the enterprise. If the hacking community has not yet discovered and capitalized on this vulnerability, the ND will not see any indications of attacks or intrusions, and may not even be aware of the hole. However, installing the patches on the routers will probably constitute at least one event, if not multiple events, and the EM will be aware of it. If the EM does not pass this information on to the ND, they will not know to keep watch for alerts reflecting an exploit of the vulnerability.



The ND needs to know that there is a vulnerability, that a patch has been developed, and also which routers have installed it and which ones have not. With this information, the ND can monitor for hits on the un-patched routers. S/he may still monitor for hits on any of the routers, but attacks against the un-patched ones may indicate that the hacker community has discovered and announced the vulnerability. The enterprise may suffer a wave of attacks.

In the above scenario, if the vulnerability was discovered by someone inside the NOSC, the ND may already be informed about it. However, many vulnerabilities are identified by organizations outside the NOSC-D, or the Air Force for that matter. This increases the


probability that the ND may not be as informed as the EM, who learns of it through the logging of events in response to the discovery of the vulnerability.

While some of the WCSS concepts may be helpful in this relationship, we currently do not see it as being a high priority since simple verbal communication of the information may accomplish what is needed. The ND may benefit from increased access to information the EM logs, such as the status of upgrades on specific pieces of equipment. But without further Cognitive Task Analysis interviews and observation, we cannot make any strong claims.

The opportunities seem to be primarily one-directional in this relationship, the EM assisting the ND. We have not learned of significant ways the ND information and operations impact the EM functions.

Network Defender and Crew Commander


One of the Crew Commander's biggest challenges is building and maintaining situation awareness about the enterprise. This includes status of equipment, status and progress of repairs, and any intrusions and attacks. Every position in the NOSC/NOSC-D has information that is essential for the CC to build an accurate picture of the above. However, the "micromanagement" issues raised earlier under the individual CC discussion hold true. Most of the information at the ND station is too detailed for the CC to build and maintain the big picture. The CC needs information formatted and presented that allows him to make overall assessments, that helps him form generalizations that he can then report to his superiors. This makes Cognitive Task Analysis research and WCSS design an attractive option. What information from the ND workstation is useful to the CC? How should this information be formatted in support of the CC global assessments of enemy activity and states of readiness?



ND-CC Research and Development Opportunities. The WCSS opportunity is classic: Support someone in building, maintaining, and, if necessary, re-acquiring situational awareness (SA). The best way to do this is to clearly articulate the critical decisions and judgments the CC needs to make, understand *why* these decisions are important, and investigate what information is available to the ND that the CC can use. Cognitive Task Analysis can fulfill this front-end analysis, and WCSS represents a useful way of presenting information to support the development of SA.

Enterprise Controller and Event Manager

The relationship between the EC and the EM shares similarities with the relationships of the ND and the EM. The EM has information that can be useful to the EC. The EM can benefit from knowledge of what kinds of events are logged, and how these are likely to show up in the enterprise health status. The EM can provide the




EC with foreknowledge of events that are likely to impact the enterprise. A WCSS that draws upon the event information that the EM maintains is a useful development for the EC.

However, there is an additional level of relationship between the EM and the EC. There should be a strong *correlation* between the status of ongoing events and the enterprise network. Problems with the enterprise should have events associated with them. Events should have corresponding impacts on the enterprise. Additional problems may exist when there is a discrepancy between the two. While we are unsure whether the correspondence will always exist (an event could exist that has no impact on the EC), the norm is probably close correspondence between logged events and enterprise health. Providing the information to the two positions that highlight the discrepancies between events and enterprise health indices would be beneficial for both positions. For example, if an event is closed but the problem is still visible to the EC, this should be a flag. In contrast, if the enterprise problem goes away but the event is still open, either an EC sensor is wrong or the EM is tracking an event that has already been resolved. A means of making these connections for the EC and EM would be beneficial. In this case we are helping the parties recognize that one or more of the expectancies (RPD) have been violated.

EC-EM Research and Development Opportunities. Like the ND-EM relationship, a Cognitive Task Analysis of the EC position could be conducted to uncover the decisions and judgments and information needs of this position. This CTA would be performed in conjunction with one of the EM, to see what *common* information the two positions could and should be sharing. The CTA results would be leveraged to create a single WCSS for both the EC and EM that displays events and their corresponding appearance in the enterprise, and make evident any discrepancies between enterprise health indicators and the status of events.

Enterprise Controller and Crew Commander.


The EC-CC relationship is very similar to the ND-CC relationship. The biggest challenge for the Crew Commander is remaining informed about the enterprise health status. The CC should be familiar with the status of the enterprise at a sufficient level of detail to report major health issues at command-level briefings. But the CC does not need the same level of information that the EC uses.



EC-CC Research and Development Opportunities. The challenge is to create a WCSS that provides the CC with an appropriate view of enterprise health issues, with the capability to drill for additional detail as needed.

Event Manager and Crew Commander

The EM-CC relationship has many of the same features as the ND-CC and EC-CC combinations, but the EM-CC relationship may be strongest of all. The CC has the responsibility to track and brief everything that is happening with the enterprise, at a level that will allow



assessments to be made regarding mission readiness and capabilities. In many respects, most of this assessment can be derived from an assessment of new and ongoing events.

The EM's role is basically to log and track these events continuously, so the CC is not overwhelmed with the task. While there are tools that help represent the status of events, we do not believe they are optimized for the level of information, and in some cases, the type of information that would be most beneficial for the CC. It is true that the CC will occasionally want to know the status of specific events. However, the CC also needs to be able to see patterns that allow him or her to make key decisions and judgments:

- NCC-X is consistently behind schedule on event resolution.
- We are having a high rate of problems with this particular model of server.
- All events requiring materials from X source will be behind schedule.
- Vendor X is handling these events and will likely finish ahead of schedule.

The CC should be able to identify significant patterns such as these.

EM-CC Research and Development Opportunities. Additional Cognitive Task Analysis interviews and observation would be needed to capture the essential information that the EM can provide in support of the CC's building and maintaining of situation awareness. We envision this display however being a member of an integrated suite that not only provides information about events and their significance, but also integrates a view of both the "friendly" (enterprise) and "enemy" (alert) situations. That is, a single WCSS could be created for the CC that bridges information from the EM, ND, and EC: patterns in events, alerts, and enterprise health issues.

AFCERT, NOSC/NOSC-D, IW FLIGHT, AND NCC: ASPECTS OF COORDINATION

In the previous section, we focused on the responsibilities and challenges associated with each of the NOSC-D positions, and the challenges associated with the coordination between the positions. This section of the report discusses the NOSC and NOSC-D interactions with other layers of the Air Force enterprise: AFCERT, IW Flight, and NCCs. It is important to acknowledge that this section of the report does not discuss detailed findings about current interactions between these different layers. We learned during our data collection trip to the NOSC-D that there is currently very limited interaction between these layers. The NOSC-D has few experiences that we could investigate. Instead, the focus of this section is a detailed research plan and rationale for investigating these interactions. A low-fidelity scenario is suggested as a research tool to study the interactions among the layers.

It appears that significant forethought and planning underlie the hypothetical, idealized configurations and interactions among AFCERT, IW Flight, NOSC, NOSC-D, and NCCs. However, our initial visits with AFCERT, ACC NOSC, 12AF NOSC-D, and multiple NCCs suggest that the actual interaction between these levels is restricted. The NOSC and NOSC-D, as relative newcomers to the AF enterprise, have not yet fully integrated with AFCERT and NCC/NCC-Ds. AFCERT and the NCCs have found ways to work together, and the NOSC/NOSC-D layer is sometimes viewed as a distraction. We also found that the 12AF NOSC-D can report little about how they interact with the IW Flight while integrated into an

AOC, beyond what the concept of operations (CONOPS) documents specify. Their deployed experience to date is largely confined to two exercises (Blue Flag, Unified Endeavor), neither of which significantly engaged the NOSC-D's full range of capabilities.

Rather than reproducing the documented interaction protocols, this section of the report presents a detailed plan of how the coordination could be studied using Cognitive Task Analysis methods. This research plan can serve as a road map for better understanding, and maximizing the potential of the tiered Air Force Enterprise Network Operations.

To facilitate an understanding of how the AF enterprise layers would actually coordinate with one another, we suggest a low-fidelity simulation exercise that investigates the interactions between these groups. Representatives from each group are brought together to work through and discuss their envisioned responsibilities, communications, and actions during a table-top exercise. This will help identify critical communication nodes between the entities, and highlight the necessary communication network and protocol to ensure that vital interactions are supported. This exercise should remain low-fidelity and table-top, so that equipment issues will not interfere with communications.

The exercise would focus on the information sharing needs, irrespective of the technology used to collect and share it. The table-top exercise would not address problems and issues associated with computers or software; that analysis would come later. Klein Associates has extensive experience in developing and implementing low-fidelity exercises aimed at identifying team decision and information-sharing requirements. These types of exercises are much easier and more cost-effective to implement when technology is not an element.

We would likely take the following approach in conducting the table-top exercise:

1. Develop and test attack scenario exercises. To develop meaningful exercises that will uncover the communication needs of all levels of the AF network security enterprise, we must first understand the basic roles of individuals at each level. Fortunately, in this effort and earlier efforts, we have interviewed individuals at NCCs, ACC NOSC, 12AF NOSC-D, and AFCERT and have that information we can leverage. We also have access to written materials (i.e., CONOPS) on NOSC, NOSC-D, and IW Flight. We would develop draft exercises based on the existing CONOPS documentation and interview data. The exercises would simulate a coordinated attack on multiple AF targets, and would provide specific information to players at each level that they would have during the attack. To ensure that these exercises would be realistic, we would involve subject-matter experts at each level of AF network security during the development process. We would also present the exercises to these experts and incorporate their feedback to ensure a final scenario that adequately represents their environments.

Creating the materials necessary to simulate a distributed attack, even in a table-top exercise, is a significant undertaking. We presently have some knowledge of the components of a distributed attack through previous research in the information warfare domain (McCloskey & Stanard, 1999; McCloskey et al., 2001). However, a more thorough analysis of the elements of a distributed attack could be investigated as a

separate, but related research effort. The final section of the report discusses a detailed distributed attack research plan.

2. Run the exercises. To conduct these exercises, we would bring representatives of each AF enterprise level to the same location.

We would include AFCERT analysts to represent AFCERT decision making. This includes personnel involved in operations (i.e., those who conduct batch analysis, real-time analysis, and incident response), and personnel from support, since they work with ASIMS (Automated Security Incident Measurement System).

For the NOSC and NOSC-D, we have found that the Crew Commander is primarily responsible for official communications outside the organization, and that the Net Defender is frequently in informal contact with AFCERT and the NCCs. The Net Defender must, above all else, ensure that an NCC in question knows that an anomalous activity occurred. Although AFCERT is in charge of incident management, they are typically far too busy to give the incident the attention it deserves. The Net Defender will often pick up on incidents that AFCERT has missed. Therefore, we would include, at a minimum, representatives for Crew Commander and Network Defender positions.

Our interviews at various NCCs suggested that many of the security administrators at the bases kept themselves compartmentalized. They monitor their own systems, responding to customer traffic tickets, and reading daily traffic and network probe reports. They do not however seem to interact much with the higher levels (NOSC, AFCERT). By having these personnel from the NCCs involved in the exercise, we will likely identify information flow that is missing. Where are the places where information should be shared, but isn't?

It would also be important to include representatives from an in-garrison NOSC, as well as a NOSC-D. Although these two organizations are very similar in the positions they staff, the deployed NOSC-D and NCC likely engage in somewhat different coordination processes. At the least, the theater where deployed forces are situated introduce unique challenges that affect their functions and coordination (as discussed earlier).

Figure 6 shows the possible players in this exercise. Overall, this exercise could easily have 2 AFCERT representatives (one from the operations side, and one from the support side), 2 NOSC representatives, 2 NOSC-D representatives and at least 2 individuals from NCCs (1 NCC, NCC-D). We would also include at least 1 representative of the IW Flight. Even if entities other than the IW Flight develop a clear picture of enemy activity, we have to ensure that this picture is effectively passed to the node that can take offensive countermeasures. By including the IW Flight, we can determine whether this interaction occurs, and if not, how it can be supported.

Each of these individuals would have a different, partial picture of a coordinated attack against multiple AF targets. We would provide different information to different individuals, based on the part of the picture they would normally have. We would provide

updated information to the individuals as an attack progresses. We would also provide additional noise data that would represent the normal traffic and distracters of their jobs. The individuals would be tasked with performing normal, routine tasks, while monitoring their individual “stations.” At any time, the individuals would have the opportunity to share information with the other participants. We would observe the information flow as well as the development (or lack of development) of the “big picture” of the attack. The greatest benefits of the exercise would be in seeing how well or poorly these entities share their information, and demonstrate how different entities will have different pieces of the big picture that must be shared to effectively identify and counter an attack.

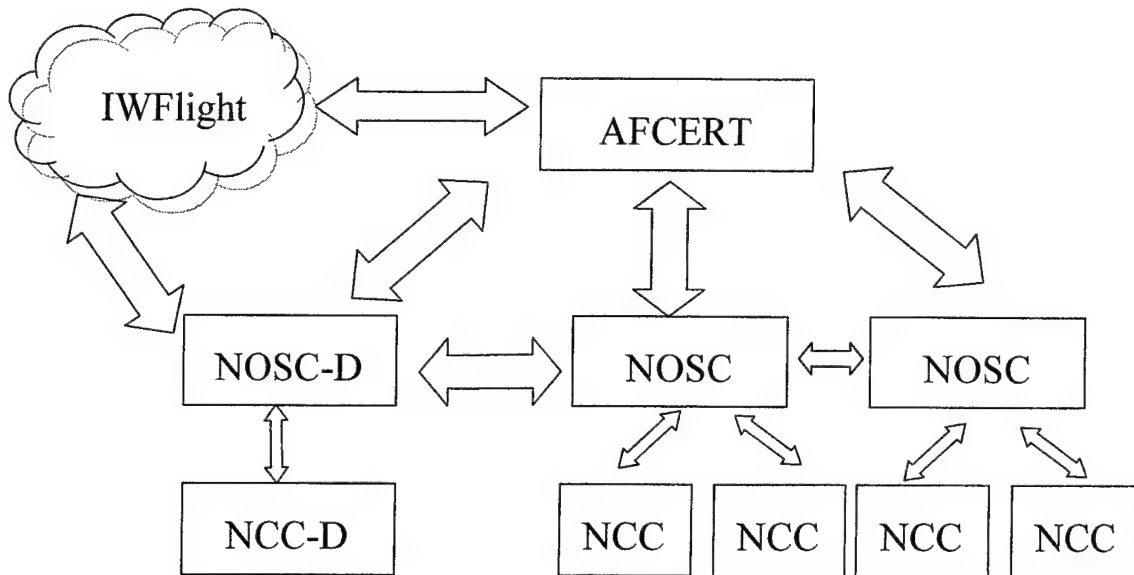


Figure 5. Air Force network security interactions.

3. **Follow-Up Interviews.** During and after the table-top exercise, we could conduct interviews with participants. We would use Team Cognitive Task Analysis methods to study the relationships between the enterprise levels and the information and decision requirements. The Wagon Wheel and Situation Awareness Calibration techniques are potential candidate knowledge elicitation strategies (Klinger et al., 2000).

SA Calibration provides insight into how individuals within a team think. At any given point in time, different team members will have different assessments of a situation. During the exercise, the exercise is stopped at a critical point and the following five questions (via a paper questionnaire) are asked of participants:

1. What is the immediate goal of your squad?
2. What are you doing to support that goal?
3. What are you worried about?
4. What is the current threat location, size, and intention?
5. What do you think this situation will look like in 20 minutes, and why?

In the Wagon Wheel technique (see Figure 7), the interviewee's name is placed in a center node. The interviewer asks the interviewee to whom s/he communicated when the team was in operation. In this case, the relevant operation is the exercise. The names are placed as nodes surrounding the interviewee's name. The next, and most time-consuming step is to probe for more information regarding the connections, or spokes, in the wagon wheel. Example probes include:

- What type of information was passed between you and X?
- From where did you receive the information you transmitted?
- What decisions does this information affect?
- What's the impact to the team if this communication line is broken?




Figure 6. Wagon Wheel method.

4. Evaluation of results / Recommendations for more effective info sharing. Based on the findings from this exercise and the participant interviews, we would generate specific recommendations for supporting the decision requirements and information exchanges between each of the AF nodes. We would specify the critical information sharing nodes, and provide guidance in terms of both organization and interface/system design (i.e., WCSS) to support the development and maintenance of a big picture of a coordinated attack. This research would also demonstrate who, if anyone in this structure, presently has the information resources to develop the big picture. Do these decision requirements and information elements exist at a single organization, or are they distributed across levels? Or do the current systems and organizational structures generate a situation such

that, even if all the information were to be collocated, it would not result in a big picture of a coordinated attack? If this is the case, we will identify what information is missing, and what practices and supports need to be implemented so the critical decisions and judgments can be made in a timely, effective manner.

Deployed Unit Attacks versus Air Force Wide Attacks

In developing and running this exercise, we must distinguish between two different types of attacks. The first is an attack against a deployed unit. Here, the adversary will be focusing their efforts specifically on the deployed assets, trying to infiltrate or otherwise damage the networked systems that are specifically used to support the deployed units. We must also consider a second scenario; that of a larger-scale attack against the Air Force. Technologically sophisticated foreign adversaries will likely understand that the US Air Force network enterprise has nodes throughout the world, and can inflict damage at a wide range of possible targets. It may behoove the adversary to target a central Air Force infrastructure, in order to limit the impact the Air Force can make any number of places. For example, even if a unit is deployed to support a specific conflict, a critical target for the adversary may be the Air Force base in the United States that schedules refueling aircraft for the strike packages. While this base is not part of the deployed unit, an attack on that base's networks can still inflict great damage or reveal very sensitive information important to the deployed unit's mission. For this research effort, we would develop separate exercises for each type of scenario (NOSC-D versus in-garrison NOSC). The information flow requirements will likely be unique for each attack scenario. The leverage points for attackers will also differ.

MODELING A DISTRIBUTED ATTACK

In the previous section, we outlined a potential research plan for investigating the coordination among levels of the Air Force enterprise network. This research plan was presented because the NOSC and NOSC-D have had few critical experiences that permit an accurate analysis of their coordinating functions with AFCERT, NCC, NCC-D, and IW Flight. Table-top exercises were suggested to study these interactions, because our preliminary analysis shows that there is presently little interaction between these nodes that can be reported, beyond the existing CONOPS documentation. We suggested two separate table-top exercises to study these interactions: one that presents an information attack targeting in-garrison forces, and other simulating attack against a deployed force.

One of the challenges associated with the creation of an exercise is adequately representing the type of attack the AF enterprise can expect. Our research to date in the information warfare domain has shown that the most formidable scenario is a massive, coordinated, distributed attack from a foreign adversary with large personnel and technical resources. Attacks from recreational hackers and crackers is not the real threat, and yet, these are the most prevalent (documented) threats the AF has directly encountered.

In this section, we report our understanding of how well the AF enterprise currently understands distributed attacks, and the limitations in the ability of the various levels to detect them. We then present a staged-model of an attack that we developed through Cognitive Task

Analysis of network attackers. While not necessarily representative of a distributed attack from a foreign adversary, the model can serve as a preliminary basis for the investigation of information warfare. We follow the description of the preliminary model with a suggested research plan for investigating the parameters of a distributed attack, as the basis for the scenarios presented in a table-top simulation exercise. The applicability of a distributed attack model also extends to the creation of WCSS for each level of the AF enterprise to support the detection and response to this critical attack type.

Detecting the Distributed Attack

If adversaries decide to launch a massive, coordinated, distributed attack on a multitude of Air Force computer networks, will the defenders be able to recognize the attack for what it is? Is there any one position that has the visibility and resources to catch early indicators of such an attack? If we assume that somewhere the visibility is clear, and the resources are available, do the security personnel have the knowledge and ability to generate an accurate mental model of what to expect?

Our exposure to AFCERT, NOSC-D, ACC NOSC, and various NCCs suggests that the answer to these questions is no. While the Air Force recognizes the potential threat posed by coordinated attacks on multiple networked targets, it appears that little attention has been given to what such attacks might look like at various levels of network security. This complicates our ability to generate a realistic table-top exercise with events and parameters that adequately mimic a distributed attack.

Consider our findings to date from preliminary interviews at all levels of the AF enterprise (McCloskey & Chrenka, 2001). One organization that has broad visibility is AFCERT. While they can see the broadest picture of overall AF network security, this comes at a cost. Our interviews found that, given their workloads and breadth of coverage, the analysts can only look at a portion of network security at each base. They must quickly move from event to event in order to keep from being overloaded. The NOSCs cover smaller numbers of bases, and they look at a much wider range of network status information. However, they had limited visibility into individual base activity, especially when it came to intrusion detection information. The ACC NOSC had to make information requests up to AFCERT when they wanted detailed activity information around a potential attack. Also, given that the NOSCs each monitor a different section of bases, they would have little knowledge of suspicious activity that occurs at bases outside their scope of responsibility. The individual NCCs closely monitor the day-to-day activity at the base level, but this appeared to be a level where the focus is on day-to-day operational capability rather than intruder detection and analysis. Inexperienced personnel are often placed in monitoring positions at the NCCs, and they have little, if any, training on recognizing patterns of attack. And obviously, personnel at this level will have the least exposure to network activity at other bases.

We found that personnel at all three levels had widely varying perceptions of how even a single enemy might attack Air Force targets. In a recent research effort (McCloskey & Chrenka, 2001), we provided security experts at AFCERT and multiple NCCs with a network topology map, and queried them on what they saw as the most likely enemy attack methods against this

configuration. Most of the AF security experts we interviewed were very technically competent, and many could clearly envision detailed routes that attackers would likely take, but no two interviewees envisioned similar paths and preferences that they would expect attackers to execute within the topology. We witnessed no consistency, even within security centers, of what actions enemies would likely take and why they would take them. Even when the interviews focused on their own networks, there was no shared understanding of how the enemy would attack.

We also found that distributed attacks cause even more challenges for these electronic battlefield warriors. One AFCERT analyst told us, "When people start in this job, they don't recognize patterns in activity over time. This comes slowly, as experience is gained. These patterns in suspicious activity need to be made more apparent." Another Air Force analyst told us, "AFCERT sticks in specific strings to search for. If the network attacker is trying to conceal their efforts (slow scans at random time intervals), AFCERT won't catch it." Clearly there is a need for the development of a shared picture of enemy activity. The only way to win a battle is to get ahead of the enemy's decision cycle. Traditional warfare has taught us this. Applied to the electronic battlefield, this means that we must first develop a coherent picture of the adversary.

When an attacker spreads his attack over an extended period of time, it is less likely to be detected, both by the human and by the computer. This task becomes even more difficult when the human operator does not have an understanding of how these distributed attacks might occur. The experienced interviewees said that they occasionally catch a distributed attack, but in order to do this effectively, you need to understand how these distributed attacks are likely to occur so that expectancies can be generated. Again, an understanding of the enemy is essential.

Given this division of responsibilities, it is unclear as to who, if anyone, would have a complete picture of a coordinated attack against multiple Air Force assets. But what picture do the NCCs, NOSCs, and AFCERT have? If a distributed attack were to occur, what would that attack look like at these levels? Taking what we have learned about network attacker decision making over the past three years, we propose to develop a picture of what a coordinated attack against distributed assets might look like, and how this attack could manifest itself at the varying levels of the Air Force network security infrastructure. This information could serve as the basis for the table-top exercises described in the previous section. The information could also serve as an important stand-alone research effort that would benefit the creation of a WCSS that supports the detection and response to a coordinated, distributed attack.

Stages of a Coordinated Attack

We have learned from previous research that attacks against network targets, whether conducted by teenage pranksters or malicious foreign entities, all seem to have distinct, yet overlapping cognitive stages (McCloskey et al., 2001) When attackers are at a particular stage, their activities should result in different indicators of attack. We have identified the following stages of attack:

Step 1: *Target selection.* Attackers need to determine a target on which to focus their efforts. In group settings, a leader will typically make the selection of what

target to attack. Often, s/he will decide this with the more advanced or trusted group members. Occasionally, a group member may suggest a particular target to attack, but the final decision generally lies with the leader. If the attacker is a foreign adversary, then general target selection may occur at a high level of command. This strategic-level decision may be bounded by specific goals. For example, the goal may be to paralyze Air Force refueling capabilities in the Middle East for 48 hours. This high-level goal may then be passed off to an operations team that would then determine which targets, if struck, would result in achievement of the high-level, strategic goal.

Step 2: External Reconnaissance. Next, the attackers will research the systems to attack, trying to identify such things as Internet addresses, types of software used, and other information that may be useful in finding holes through which to enter the systems.

Step 3: Attack. Once this identifying information is found, attackers may apply preexisting exploit tools or try standard techniques to gain access into the system. Occasionally, attackers will engage in "social engineering" to attempt to subtly draw information out from system administrators or other target personnel. For example, an attacker may telephone a security administrator, claiming to be an irate customer who is having trouble accessing their system. They will then try to gain as much information as possible from the administrator, who may be trusting and try to assist the caller. Within the AF enterprise, they could call a NOSC/NOSC-D Event Manager under the guise of NCC personnel, and collect information about target systems.

Step 4: Internal Activity. Once into the system, skilled attackers will typically check system log files to ensure that they cannot be identified from the recorded information. They may then identify the internal software that is running on the network and then investigate known weaknesses of that software (particularly the network operating system). Attackers may also search for other systems to which the hacked system is connected. This is done so that attackers can continue to gain more access into additional systems. Finally, some attackers set up "sniffers" on the system. These sniffers are data recorders that keep track of conversations and interactions that occur over the network. The sniffers dump raw data into a log file. Although these log files contain vast amounts of useless information, they also record information about passwords of legitimate users, and the connections legitimate users make between systems. Network attackers have automated programs that identify, in these vast log files, this vital information. Again, this information can be used to provide new access to different systems for the hacking group.

Step 5: Post-hack Activity. Once the attack has been conducted and the attacker has left the scene, there is often still work to be done. Depending on the goals of the attacker, different activities may be observed. If the attackers were looking to impress, or coordinate, with a small group of colleagues, they will inform these

few individuals and provide them with information on how to access the hacked system for themselves. If the attackers are looking for wide-spread recognition, they may alert the media, or make postings to multiple sites about the attack. If the attackers are foreign adversaries, their next actions may be at a strategic level. If, for example, the attack resulted in an inability for the Air Force to refuel its fighters in the Middle East, the foreign adversaries may conduct physical attacks on targets the US aircraft are temporarily unable to protect, due to their inability to refuel.

Although these stages appear to form a linear process, the stages of a cyber-attack are not necessarily sequential. For example, if external reconnaissance on a selected target uncovers that the target is guarded too tightly, attackers may go back into the target selection stage to identify a more suitable target to attack.

It is important to note that, for our earlier research, we focused primarily on activities of individuals, or socially-oriented hacking groups. We need to research the nature of distributed, coordinated attacks from foreign adversaries. While we can learn a great deal by drawing connections between traditional hacker decision making and foreign adversary decision making, we cannot treat them as the same. We need to understand the specific decision making that occurs when foreign threats plan and launch large-scale attacks against US Air Force targets.

While it would be impractical to propose to study these foreign adversaries directly, we can offer a different approach. We could conduct a Cognitive Task Analysis with multiple, unconnected "red teams" in the military and civilian domains. We would pose a simulated challenge to these teams to develop a plan for conducting a large-scale attack against the Air Force enterprise. As they plan, we would conduct interviews to determine how they would select multiple targets, how they would perform their reconnaissance on their targets, and how they would coordinate such a large-scale attack.

This approach would provide data that would allow us to develop multiple models of distributed network attacks against the Air Force. We would have a picture of how targets could be selected, reconnoitered, and attacked. We would have a better understanding of the coordination issues that would arise when a distributed attack is conducted. While this picture would be based on a series of red teams, rather than actual adversaries, it would still be accurate enough to allow the formation of expectancies. We will likely find that certain things must occur in a distributed attack, regardless of the attackers.

The next step would be to define just how such attacks would appear to: 1) individual NCCs that are hit, 2) NCCs that aren't hit, 3) NOSCs that are overseeing the attacked NCCs, 4) NOSCs that are overseeing non-targeted NCCs, and 5) AFCERT. This is critical for the development of a table-top exercise having players from each level. Each entity will likely have a different view into the coordinated attack. We need to understand which pieces of the puzzle are found at each of these locations. What indicators will be visible at one location, but not at the others? As we study this, we also need to understand current practices. What are personnel currently doing with the information they receive on attacks? What decisions are being made and what information is being passed? Is there anyone who is constructing the big picture?

Looking back on the stages of a network attack that we identified in previous research, we propose to look at each stage individually, and determine first, how adversaries to US Air Force interests could conduct these stages in a coordinated fashion, against multiple targets. We would want to identify the key leverage points of the Air Force network security enterprise from an enemy's point of view. What are key combinations of targets that foreign powers could strike in order to inflict extensive damage? How would they select these combinations of targets? How would they reconnoiter them, and how would they coordinate the actual attacks, including sequencing of targets? Next we would determine the indicators of these stages at different locations. For example, when adversaries are reconnoitering targets for an upcoming strike, what visibility would the different NOSC's have that this activity is occurring? What would AFCERT see? What activity could be observed, but is not currently monitored?

This research would first determine where the individual elements of the big picture in detecting distributed attacks exist. It would also indicate whether any one entity has enough visibility into enough of these elements to develop and maintain the big picture of a distributed attack. This information could be packaged as the content for the table-top exercise described in the earlier section. It could also serve as the basis for the evaluation of the existing software various levels of the AF enterprise, and positions with the NOSC/NOSC-D use, in terms of their ability to detect a distributed attack and generate an accurate picture of it.

A Distributed Attack Against Multiple Air Force Targets

An overall goal of the research effort described here would be to document not only different, likely adversarial courses of action against the broad, distributed Air Force security enterprise, but also to show how each stage of these attacks would manifest itself at the NOSC-D and other Air Force security nodes. This would tell us where the critical information to detect and respond to attacks resides as well as what communication channels need to be reinforced. Following this research, we should have detailed, specific information for each of the cells in Table 2. Note that we have included the type of information that would populate some cells.

To study the events and parameters of a distributed attack, we would pursue the following:

1. Experienced Network Attacker Interviews. Using our existing established connectivity to the hacker underground (truly expert hackers who don't publicize), we would conduct a brief series of interviews with highly experienced hackers who have worked in teams to determine how they would select, reconnoiter, and attack multiple targets using distributed tactics to inflict the most damage and avoid detection. We would focus our interviews on the specific strategies that hacker teams use when faced with massive targets, along with what they would perceive as potential indicators of their activities. By understanding how these hackers attack large-scale targets, we should get a sense of how foreign adversaries might identify leverage points in the Air Force network security perimeter. An important note is that we would create methods to collect this information from hackers, without inadvertently revealing information about specific Air Force targets that the hackers might try to attack later.

Table 2.

Views into a Distributed Attack at Different Locations

Location/ Stage of Attack	Targeted NCCs or NCC-Ds	Non-targeted NCCs	IW Flight	NOSCs or NOSC-Ds covering targeted NCCs	NOSCs covering non-targeted NCCs	AFCERT
Target Selection	?	?	?	Intel from AFCERT?	?	Possible Intel on enemy intentions
Reconnaissance	Possible increased ping/scan activity at specific NCC/ NCC-Ds	?	?	?	?	?
Attack	Anomalous activity within log files, various detected attack signatures	?	?	Possibly ID alarms; patterns of alerts across NCCs	?	Alarms patterns within specific NOSCs/ NOSC-Ds
Internal Activity	Anomalous behavior from user(s)	?	?	?	?	?
Post-Attack Activity	Degraded systems/ missing, altered info	?	?	Possible patterns of enterprise health reports from affected NCCs	?	Possible patterns of enterprise health reports from affected NOSC/ NOSC-Ds

2. Red Team Interviews. We would then conduct a larger series of cognitive simulation interviews with established "red teams" in both academia and industry to understand how they would plan and launch distributed attacks. We would focus on, among other things, how they identify leverage points for attack when targets have several possible individual strike points (such as the US Air Force), and how they would exploit these leverage points. We would pose hypothetical situations (e.g., USAF has been deployed to a certain location for a certain mission) and interview these experts to determine how they would identify the computing/networking resources that are deployed, and how they would then identify the weaknesses in these resources and the most damaging and likely attack methodologies they would employ.
3. AF Enterprise Interviews. In conjunction with the red team interviews, we would also conduct Cognitive Task Analysis interviews with key personnel from: IW Flight, NOSC-D, AFCERT, other NOSCs, and NCCs. We would focus these interviews on specific systems and information requirements these personnel currently utilize, and how this information is used. We would also focus part of the interviews on situation awareness in terms of large-scale attacks. Who, if anyone, has the big picture?
4. Data Analysis. In our analysis of these data, we would identify and represent several likely courses of action that an adversary might take when attacking Air Force networked systems. This information is critical to whoever is, or will be, responsible for maintaining the big picture of Air Force network security. We would also identify how these attacks would look to the different security nodes as they are currently configured and staffed. We would identify which pieces of the puzzle are at which locations, and which pieces are not being picked up at all. Finally, we would develop recommendations for information-sharing channels that would support development of the big picture.

SUMMARY

These are exciting times for researchers interested in information security, human-computer interaction, and organizational research issues. The Air Force has created a basic foundation for securing and leveraging their information resources in battle, including the development of a three-tiered organization for in-garrison defense, and an organizational model of deployed information operations. AFCERT and NCCs are currently staffed, technically equipped, and perform their intended functions. The NOSC and NOSC-D have more recently been stood-up as a middle layer in the AF enterprise. The NOSC and NOSC-D in particular, have not yet fully formalized their role and technical proficiency within the AF enterprise. This means there is ample opportunity to identify NOSC and NOSC-D organizational and technical needs, and develop workable solutions that will make their transition into the enterprise a successful one.

This report presented a preliminary analysis of the challenges that currently face individual positions within the NOSC and NOSC-D, as well as the challenging relationships between these positions. Cognitive Task Analysis tools are suggested as a user needs analysis tool that can more fully explore these challenges and develop guidelines for interface designers. Work-Centered Support Systems is a technology solution for these challenges that seeks not to

reinvent the wheel, but build upon the commercial information technologies the Air Force is already using within the NOSC/NOSC-D. WCSS can draw together information from the separate, stove-piped information systems the individual positions utilize to increase shared situation awareness within the NOSC/NOSC-D team.

The Air Force enterprise will also benefit by low-tech research solutions, such as the table-top exercises suggested here, to study the communication and coordination of the various enterprise levels under a simulated, distributed information attack. Formal exercises to date have not fully tasked the NOSC-D, and take enormous technical and human resources to organize. Lower-fidelity exercises can be created with fewer resources that target the specific coordination and information sharing demands of AFCERT, NOSC/NOSC-D, NCC/NCC-D, and IW Flight.

REFERENCES

Beach, L. R. (1993). Image theory: Personal and organizational decisions. In G. A. Klein, J. Orasanu, R. Calderwood, & C. E. Zsombok (Eds.), Decision making in action: Models and methods (pp. 148-157). Norwood, NJ: Ablex.

Bradford, J. (2000a). AFED suggestions paper for the AFED graphical user interface (Final Report prepared under contract GS35F0764J for Air Force Research Laboratory, Human Effectiveness Directorate (AFRL/HECA), Wright-Patterson AFB, OH). Dayton, OH: Adroit Systems, Inc.

Bradford, J. (2000b). NOSC interim report for the AFED graphical user interface (Interim Report prepared under contract GS35F0764J for Air Force Research Laboratory, Human Effectiveness Directorate (AFRL/HECA), Wright-Patterson AFB, OH). Dayton, OH: Adroit Systems, Inc.

Bradford, J. (2000c). System requirements specification (SyRS) for the AFED graphical user interface (Interim Report prepared under contract GS35F0764J for Air Force Research Laboratory, Human Effectiveness Directorate (AFRL/HECA), Wright-Patterson AFB, OH). Dayton, OH: Adroit Systems, Inc.

Calderwood, R., Crandall, B. W., & Klein, G. A. (1987). Expert and novice fireground command decisions (Final Report under contract MDA903-85-C-0327 for the U.S. Army Research Institute, Alexandria, VA). Fairborn, OH: Klein Associates Inc.

Crandall, B., & Calderwood, R. (1989). Clinical assessment skills of experienced neonatal intensive care nurses (Contract 1 R43 NR0191101 for The National Center for Nursing, NIH). Fairborn, OH: Klein Associates Inc.

Eggleston, R. G., Young, M. J., & Whitaker, R. D. (2000). Work-centered support system technology: A new interface client technology for the battlespace infosphere. Proceedings of IEEE, NAECON. Dayton, OH.

Hoffman, R. R., Crandall, B. W., & Shadbolt, N. R. (1998). Use of the critical decision method to elicit expert knowledge: A case study in cognitive task analysis methodology. Human Factors, 40(2), 254-276.

Klein, G. (1998). Sources of power: How people make decisions. Cambridge, MA: MIT Press.

Klinger, D. W., Phillips, J., & Thordsen, M. (2000). Handbook for Team Cognitive Task Analysis (Report prepared under Aptima Contract #0009). Fairborn, OH: Klein Associates Inc.

McCloskey, M. J. (2001). Decision making in the information operations domain (Final Report Aptima Subcontract Number 0045-1065). Fairborn, OH: Klein Associates Inc.

McCloskey, M. J., & Chrenka, J. E. (2001). Anticipating cyber-threats: A cognitive analysis of network attack, detection, and threat perception (Final Report prepared under contract F30602-00-C-0131 for Air Force Research Laboratory, Rome, NY). Fairborn, OH: Klein Associates Inc.

McCloskey, M. J., & Stanard, T. (1999). A red team analysis of the electronic battlefield: A cognitive approach to understanding how hackers work in groups. Proceedings of the Human Factors and Ergonomics Society 43rd Annual Meeting, 1, 179-183.

McCloskey, M. J., Stanard, T. W., & Armstrong, A. A. (2001). Cognitive analysis of the information security domain (Final Report prepared under contract F41624-98-C-6002 for Air Force Research Laboratory, Rome, NY). Fairborn, OH: Klein Associates Inc.

Montgomery, H. (1993). The search for a dominance structure in decision making: Examining the evidence. In G. A. Klein, J. Orasanu, R. Calderwood, & C. E. Zsombok (Eds.), Decision making in action: Models and methods (pp. 182-187). Norwood, NJ: Ablex.

Pennington, N., & Hastie, R. (1993). A theory of explanation-based decision making. In G. A. Klein, J. Orasanu, R. Calderwood, & C. E. Zsombok (Eds.), Decision making in action: Models and methods (pp. 188-201). Norwood, NJ: Ablex.

Pliske, R. M., Klinger, D., Hutton, R., Crandall, B., Knight, B., & Klein, G. (1997). Understanding skilled weather forecasting: Implications for training and the design of forecasting tools (Technical Report No. AL/HR-CR-1997-0003 for the Air Force Material Command, Armstrong Laboratory, Human Resources Directorate Brooks AFB, TX). Fairborn, OH: Klein Associates Inc.

Rasmussen, J. (1993, July/August). Diagnostic reasoning in action. IEEE Transactions on Systems, Man, and Cybernetics, 23(4), 981-991.

Sweeney, M. J. (2000). Loveletter worm/virus lessons learned data. Langley AFB, VA: ACC/SCN.

Young, M., Eggleston, R. G., & Whitaker, R. D. (2000). Direct manipulation interface techniques for users interacting with software agents. Paper presented at the NATO/TRO Symposium on Usability of Information in Battle Management Operations, sponsored by Human Factors and Medicine (FHM), Oslo, Norway.

Zsombok, C. E., Klein, G., Kyne, M. M., & Klinger, D. W. (1992). Advanced team decision making: A developmental model (Contract MDA903-90-C-0117 for U.S. Army Research Institute for the Behavioral and Social Sciences). Fairborn, OH: Klein Associates.